

Low delay and low loss packet switch.

The present invention relates to a Telecommunication switching arrangement for switching digital data contained in data packets provided with a packet header, which contains a first and a second index representative of a first and a second characteristic of the data packet, the arrangement comprising a crosspoint switch for switching packets from a plurality of input lines of the switch to an output line of the switch, and a sequencer in the crosspoint switch for determining a sequence in which the data packets are to be switched to the output line, and a storage for storing the data packets.

A telecommunication switching arrangement according to the preamble is known from US patent USA5,390,176. This telecommunication switching arrangement determines in which sequence the incoming data packets are to be switched to the output of the switching arrangement. The data packet contains a bit in the header that indicates the type of package. There are two types of packets: lowloss packets and low-delay packets. The switching arrangement contains a buffer to temporarily store packets for which there is insufficient transmission capacity available. Incoming packets are evaluated by the switching arrangement. Low-loss packets always have storage priority over low-delay packets. Low-delay packets are only stored when sufficient storage capacity is available. This results in a switching arrangement where low-delay packets are processed more quickly than low-loss packets, while low-loss packets are less likely to be lost than low-delay packets.

In view of recent developments towards switching arrangements that provide Quality of Service it is insufficient to provide only two service classes for the data packets. Several aspects of the data packet must be considered, as well as several priority levels within one aspect.

A disadvantage of the telecommunication switching arrangement according to US-A 5,390,176 is that the switching arrangement is incapable of determining the sequence in which data packets with several priority levels per aspect are to be switched to the output line.

An object of the present invention is to provide a telecommunication switching arrangement according to the preamble that can determine the sequence in which the data packets with several priority levels per aspect are to be switched to the output line.

To achieve said object the present invention is characterized in that the sequencer is operative to determine the sequence in which the data packets are to be switched to the output line based on a result of a mathematical operation on the first index and the second index.

The data packet contains two indices in the header. Each index represents an aspect of the data packet. The index has a range, and the value of the index represents the relative priority of the data packet within the aspect. Both indices together define a service class to which the data packet belongs.

By combining the first and second index of the data packet by means of a mathematical operation into one result, it is possible to create groups of different service classes, where the result indicates the group to which the data packet belongs. Depending on the mathematical operation the various aspects of the data packet can be given a certain weight in the result of the mathematical operation. The sequencer compares the result of the mathematical operation of several data packets thus establishing clear relative priorities between the data packets. Based on these relative priorities between the data packets the sequencer can determine the sequence in which the data packets are to be switched to the output line.

An embodiment of the present invention is characterized in that the mathematical operation is the addition of the first index and the second index.

The first and second index are added, and the sequencer uses the result of the addition to establish the priority between the data packets. If a lower value of an index is equivalent to a higher priority, the result of the addition of two indices with low values is again a low value, representing a data packet with high priority. The result of the addition of two indices with high values is a high value representing a data packet with a low priority.

A further embodiment of the present invention is characterized that the mathematical operation is the multiplication of the first index and the second index.

The first and second index are multiplied, and the sequencer uses the result of the addition to establish the priority between the data packets. If a lower value of an index is equivalent to a higher priority, the result of the multiplication of two indices with low values is again a low value, representing a data packet with high priority. The result of the

multiplication of two indices with high values is a high value representing a data packet with a low priority.

A further embodiment of the present invention is characterized that the sequence for packets with an equal result of the mathematical operation of the first and second index is based on the index that is representative of a predetermined packet characteristic of the data packet.

In the event that the results of the mathematical operation of data packets are equal the sequencer uses a predetermined index in the header of the data packet to establish the relative priority of the data packets. The predetermined index is this way effectively being used twice in establishing relative priority, through the mathematical operation and by using it as a sub-index. This allows the sequencer to tailor the sequence to the availability of particular resources, that are related to the predetermined index. If insufficient resources are available the sequencer can determine which data packets will be handled and which data packet can be delayed by storing the data packet in the storage or even discarded.

The present invention will now be explained with reference to the drawings.

Figure 1 shows a telecommunication switching arrangement according to the invention

Figure 2 shows the establishment of service classes when the indices are added.

Figure 3 shows the division of the storage to accommodate the groups.

Figure 4 shows a section of the storage that is used as a logical queue for a group.

Figure 5 shows a transmission sequence in which each group has an assigned transmission capacity.

Figure 6 shows the structure of the data packet

The switching arrangement 1 in figure 1 comprises a crosspoint switch 3 that is connected to a storage 5. Incoming data packets containing a header arrive at an input 7 and an input 9. The crosspoint switch 3 comprises a sequencer that extracts a first and a second index from the header of each data packet and performs a mathematical operation on the first and second index which results in a value which indicates the priority of the data

packet. Based on a result of the mathematical operation the sequencer determines if and where the data packet should be placed in the storage 5. When transmission capacity is available at an output 11 the sequencer 3 fetches a data packet with the highest priority from the storage 5 and transmits the data packet via the output 11.

Figure 2 shows the establishment of groups in the case that the indices are added.

The first index L is a loss index that indicates the importance of the data packet not being lost. The second index D is a delay index that indicates the importance of the data packet not being delayed. Although figure 2 is based on addition as the mathematical operation, other mathematical operations are possible in combination with a corresponding range of values for the indices as long as the result of the mathematical operation consistently assigns the data packet to a group having the proper priority.

In figure 2 a low value of an index indicates that the correspondent aspect is highly important while a low value indicates a low importance. The value of an index is effectively a ranking of the data packets for the corresponding aspect.

The header of a data packet can contain for example the values $L=1$ and $D=2$ indicating that it is very important that the data packet will not be lost while it is less important that the data packet is not delayed. The precise combination of one L value and one D value is called a service class. Data packets having identical L and D values in their headers are members of the same service class. The notation used for a service class is $xL\&yD$ where x is the L value and y is the D value.

This makes it possible to create an unlimited number of service classes. In this example the range of x is $1 \dots 3$ and the range of y is $1 \dots 4$, resulting in 12 service classes of data packets.

By adding x and y, the L value and the D value, an index sum σ is obtained that can be used to determine the priority of the data packet. The index sum σ is called a group. A $3L\&1D$ data packet, a $2L\&2D$ data packet, and a $1L\&3D$ data packet are all members of the same group with $\sigma = 4$ and hence have the same priority. Each group gets a share of the transmission capacity.

The use of 12 service classes in this example leads to 6 groups: $\sigma = 2 \dots 7$.

Figure 3 shows the division of the storage to accommodate the groups.

To each group a section of the storage 5 is assigned that functions as a logical queue for the data packets that belong to that group σ . The figure shows logical queues for the groups $\sigma = 2, 4, 5, 6$.

The sections can be fixed or variable, allowing the sequencer 3 to optimize the use of the storage 5 by increasing or decreasing the logical queues depending on the size of the storage 5 and the storage requirements of each group. In figure 3 no data packets of the group with index-sum $\sigma = 3$ are to be stored and hence no section of the storage 5 was assigned to this group.

Figure 4 shows a section of the storage that is used as a logical queue for a group.

In this example the logical queue of the group with index-sum $\sigma = 4$ is shown. The data packets are sorted within the logical queue so that data packets with a lower value of the delay index (D-value) are first to be fetched from the storage section. When new data packets arrive that are member of a group but have a lower value of the delay index than some data packets already in the logical queue the data packets is placed in front of those data packets already in the logical queue. This results in a sorting of the data packets in the logical queue based on their D-value, ensuring that data packets with a low D-value in a group are transmitted first. Data packets in a group that have an identical D-value are placed in chronological order to ensure the lowest delay possible for each data packet.

When data packets of group σ are being fetched from the storage the sequencer will give priority to the service class (of that group) with the lowest D-value. A threshold Δ_σ is defined to which the number of data packets of each service class is compared. The sequencer checks whether any of the service classes contain a number of data packets with a lower L-Value than the one currently being fetched that exceeds the threshold Δ_σ . If so the sequencer switches, after completion of the current data packet, to the service class in the current group that exceeds the threshold Δ_σ most. If none of the service classes exceed the threshold Δ_σ the sequencer continues fetching data packets from the current service class.

The loss index (L-value) determines, when the storage is full, which packets will be stored for further processing and which ones will be discarded. When a data packet of service class 2L&2D needs to be stored then data packets of service class 3L&4d, 3L&3D, 3L&2D and 3L&1D (in this order) will be purged from the storage to make room for the data

packet with service class 2L&2D. If there are no data packets that can be purged the new data packet will not be stored and will be lost.

The measures of the present invention thus ensure low loss and low delay for the data packets to be transmitted by effectively handling storage priority and putting the data packets in an optimal sequence for transmission.

Figure 5 shows a transmission sequence in which each group has an assigned transmission capacity. The transmission capacity can be assigned according to the Minimum Fair Queuing principle. With Minimum Fair Queuing each group has its own logical queue and out of each logical queue data packets will be fetched until the logical queue is empty or a predetermined maximum of data packets has been fetched from the logical queue, whichever comes first. The predetermined maximum is based on reservations made by protocols such as RSVP. Minimum Fair Queuing prevents starvation, that is assuming that the total reserved capacity is not higher than the available capacity, each logical queue gets at least what it has reserved. Assuming that there are no data packets in group $\sigma = 3$ to be transmitted, no capacity has been reserved for this group.

Figure 6 shows the structure of the data packet.

The data packet 23 contains a header 21 and a data section 27. The header contains the first index 25 and the second index 26.